SYSTEM AND METHOD FOR TRACKING END USERS IN A LOADBALANCING ENVIRONMENT

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to the field of communications and more particularly to a system and method for tracking end users in a loadbalancing environment.

BACKGROUND OF THE INVENTION

Data networking architectures have grown increasingly complex. In addition, the increase in clients or end users wishing to communicate in a data networking environment has caused many networking configurations and systems to respond by adding elements to accommodate the increase in networking traffic. Communication tunnels may be used in order to establish or to gain access to a network, whereby an end user or an object is tracked through a central location and provides a communication platform for the end user for that communication session.

As the subscriber base of end users increases, proper routing and efficient management of communication sessions and data flows becomes even more critical. Several components are generally implicated in loadbalancing scenarios and, in certain instances, many of these components are delegated a disproportionate number of duties. This overburdening may decrease throughput and inhibit the flow of network traffic, causing congestion or bottlenecks in the system. Additionally, the overwhelming burden on a single element in the communications flow may decrease bandwidth capabilities, as the overtaxed component is forced to work its way through each of its tasks and duties before being able to accommodate additional communications tunnels or end users. Strains on a central location not only slow network traffic, but such a reliance on a single central node may create severe problems in situations where the central node fails or is otherwise unable to perform all of its assigned duties.

SUMMARY OF THE INVENTION

From the foregoing, it may be appreciated by those skilled in the art that a need has arisen for an improved communications approach that provides for a reduction in the burden placed on a loadbalancer associated with communications between multiple end points or nodes. In accordance with one embodiment of the present invention, a system and method for tracking an end user in a network environment are provided that greatly reduce disadvantages and problems associated with conventional loadbalancing techniques.

According to one embodiment of the present invention, there is provided a method for tracking information in a network environment that includes receiving a packet included within a request that is associated with an end user and communicating the packet to a selected one of a plurality of gateways. An object may be built that correlates an internet protocol (IP) address associated with the end user to the selected gateway such that the object may be used to direct subsequently received packets associated with the end user to the selected gateway. Subsequently received packets may be directed based on end user IP information included in the subsequently received packets.

Certain embodiments of the present invention may provide a number of technical advantages. For example, according to one embodiment of the present invention a communications approach is provided that allows a loadbalancer to track end user sessions based solely on their IP addresses, instead of tracking end user sessions on a per-flow basis. This reduction in responsibility for the loadbalancer operates to increase throughput, as

two points or nodes may efficiently communicate while the loadbalancers involved need only inspect the destination and source IP addresses. This may also reduce the number of central processing unit (CPU) cycles that may be intensive and that may require additional work to be performed by the loadbalancer. Accordingly, the loadbalancer may be relegated to simplified information transfer and object building that may be used to correlate an end user with a selected gateway. The removal of per-flow tracking by the loadbalancer may further alleviate its responsibilities such that greater bandwidth may be accommodated for in the network.

Yet another technical advantage associated with one embodiment of the present invention is also a result of IP address-based tracking. The decreased reliance on flow tracking by the loadbalancer conserves memory resources. This is because the loadbalancer does not need to store flow information to load balance user sessions in the network. The loadbalancer may reduce identification of the end user to a 4-byte IP address (e.g. using internet protocol version four). This procedure may better allocate network resources because per-flow information is no longer needed to be maintained by the loadbalancer. In addition, the decreased dependency on the loadbalancer allows for improved failover characteristics such that, if the loadbalancer should become dysfunctional or non-operational, associated communication sessions may not necessarily be lost. Furthermore, redundancy problems are less challenging because of the reduction of duties assigned to the loadbalancer.

Still another technical advantage associated with one embodiment of the present invention relates to the configuration of the loadbalancer. The configuration employed is simplified in accordance with the teachings of the present invention such that less hardware, memory, table space, and/or application specific integrated circuits (ASICS) are needed to effectuate the operations of the loadbalancer. This in turn may allow for better space allocations, which may allow additional end users to be accommodated, as net aggregate capacity/throughput is thereby increased. This may also result in greater throughput being provided to each of the corresponding gateways. In essence, layer-three path enforcement may offer more scalable loadbalancing with higher throughput and total host capacity. This efficiency improvement is applicable to a number of communication environments and may further provide a scaling ability as end user subscription bases grow. Certain embodiments of the present invention may enjoy some, all, or none of these advantages. Other technical advantages may be readily apparent to one skilled in the art from the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

To provide a more complete understanding of the present invention and features and advantages thereof, reference is made to the following description, taken in conjunction with the accompanying figures, wherein like reference numerals represent like parts, in which:

FIGURE 1 is a simplified block diagram of a communications system for tracking an end user in a loadbalancing environment in accordance with one embodiment of the present invention; and

FIGURE 2 is a simplified timing diagram illustrating an example interaction between multiple elements that may be included in the communication system.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS OF THE INVENTION

FIGURE 1 is a simplified block diagram of a communication system 10 for tracking an end user 12 in a network environment. Communication system 10 may include end user 12 (which may also be referred to herein as 'Host A'), a set of loadbalancers 14 and 18 (represented as loadbalancer #1 and loadbalancer #2), a set of gateways 20 and 22 (represented as gateway #1 and gateway #2), and a set of servers 28 and 30 (represented as server #1 and server #2). Note that the designation of numerous elements included within communication system 10 as either '#1' or '#2' is arbitrary and has been provided for purposes of offering a clear discussion of some of the teachings of the present invention. These designations should not be construed in any way to limit the broad scope and considerable adaptability of the present invention in any number of communications environments.

FIGURE 1 may be generally configured or arranged to represent a 2.5G communication architecture applicable to a Global System for Mobile (GSM) environment in accordance with a particular embodiment of the present invention. However, the 2.5G architecture is offered for purposes of example only and may alternatively be substituted with any suitable networking protocol or arrangement that provides a communicative platform for communication system 10. For example, communication system 10 may cooperate with any version of a gateway packet radio service (GPRS) tunneling protocol (GTP) that includes an ability to track an internet protocol (IP) address. Communication system 10 may also be inclusive

of first generation or 3G architectures that provide similar features for identifying an IP address. Such applications for communication system 10 may extend beyond the mobile space, such as, for example, in firewall architectures.

In accordance with the teachings of the present invention, communication system 10 operates to alleviate the responsibilities associated with loadbalancers 14 and 18 in providing efficient communication using per-end user IP address tracking. In other scenarios, gateway loadbalancing is generally accomplished using flow-based loadbalancers that track each layer-four flow generated by a host. The loadbalancers can associate each of these flows to a gateway, which is tracking the host's traffic. Flow tracking alone is not sufficient to enforce host-to-gateway binding, so an additional table is generally created to map host addresses to the appropriate gateway. The combination of flow tracking and host tracking functionally enforces the gateway's requirements of binding a host's traffic to a particular gateway.

In contrast to flow tracking, communication system 10 offers enhanced path enforcement in each of loadbalancers 14 and 18 by tracking only the host's IP address (layer-three source or destination, depending on which side of the gateway the loadbalancer is positioned and the direction of the flow of the packet, relative to the gateways). Thus, end user 12 information may be tracked so that packets from the host side that are received at loadbalancer 14 (where the source information is equal to the IP address of Host A), may be directed to gateway (#1) 20. From the origin side, where the

destination IP address is equal to Host A, the packets may be routed to gateway (#1) 20. Thus, only 8-bytes of information is needed in order to properly direct packets (4-bytes for the IP address of the client and 4-bytes for the IP address of the associated gateway). Fewer elements are being tracked in the loadbalancing process. Thus, the amount of information that needs to be stored and searched within loadbalancers 14 or 18 is minimized, thereby increasing the capacity of the loadbalancer and, ultimately, the throughput of the gateway cluster.

Communication system 10 allows each loadbalancer 14 and 18 to load balance end user data sessions based on the host IP address, instead of load balancing end user data sessions on a per-flow basis. This reduction in responsibility for loadbalancers 14 and 18 increases throughput because multiple points or nodes within the network may efficiently communicate while loadbalancers 14 and 18 need only inspect the destination and source IP addresses. This may also reduce the number of central processing unit (CPU) cycles. Accordingly, loadbalancers 14 and 18 may perform simplified information transfer and object building that may be used to correlate end user 12 with a selected gateway 20 or 22. The removal of per-flow tracking by loadbalancers 14 and 18 may further alleviate their responsibilities such that greater data throughput may be achieved in the network.

The decreased tracking performed by loadbalancers 14 and 18 also conserves memory resources. This is because loadbalancers 14 and 18 do not need to look up flow keys or perform other equivalent operations. Loadbalancers 14 and 18 may simply load balance a user data flow based on a search for a host IP address in the host IP address

table. Note that IP version four (IPv4) addresses are four bytes and IP version six (IPv6) addresses are sixteen bytes. With respect to IPv6, storage may be reduced to sixteen bytes. The advantage in implementing communication system 10 generally remains across various layer three protocols/versions (e.g. IPX). This procedure may better allocate network resources because per-flow information no longer needs to be maintained by loadbalancers 14 and 18. Additionally, the configuration of loadbalancers 14 and 18 is also advantageous because of its simplified architecture. Loadbalancers 14 and 18 may be designed such that less hardware, memory, table space, and application specific integrated circuits (ASICS) are needed to effectuate their operations. This increase in space allocation may allow for additional end users to be accommodated, as an increase in net aggregate capacity/throughput is thereby provided. This may also result in greater throughput being provided to the corresponding gateways 20 and 22, which are not burdened by sophisticated tracking operations. Layer-three path enforcement may offer more scalable loadbalancers 14 and 18 and gateway clusters 20 and 22. This efficiency strategy is adaptable to a number of communications environments, and may further provide the ability to scale a given system as end user bases grow. The combination of flow-less loadbalancing and end-user IP address stickying accomplishes this function without tracking cumbersome communication flows.

End user 12 is a client or a customer wishing to initiate a communication in communication system 10. End user 12 may be inclusive of devices used to initiate a communication, such as a computer, a personal digital

assistant (PDA), a laptop or an electronic notebook, a telephone, a mobile station, or any other device, component, element, or object capable of initiating voice or data exchanges within communication system 10. End user 12 may also be inclusive of a suitable interface to the human user, such as a microphone, a display, a keyboard, or other terminal equipment (such as for example an interface to a personal computer or to a facsimile machine in cases where end user 12 is used as a modem). End user 12 may also be any device that seeks to initiate a communication on behalf of another entity or element, such as a program, a database, or any other component, device, element, or object capable of initiating a voice or a data exchange within communication system 10. Data, as used herein in this document, refers to any type of numeric, voice, video, or script data, or any type of source or object code, or any other suitable information in any appropriate format that may be communicated from one point to another.

End user 12 may be coupled to a suitable IP network in order to access loadbalancer 14. The IP network represents a series of points or nodes of interconnected communication paths for receiving and transmitting packets of information that propagate through communication system 10. The IP network may offer a communicative interface between end user 12 and selected locations within the network. The IP network may be any local area network (LAN), wireless local area network (WLAN), metropolitan area network (MAN), virtual private network (VPN), wide area network (WAN), or any other appropriate architecture or system that facilitates communications in a network environment. The IP network

may implement a transmission control protocol (TCP)/IP or a user datagram protocol (UDP)/IP communication language protocol in particular embodiments of the present invention. However, the IP network may alternatively implement any other suitable communication protocol for transmitting and receiving data packets within communication system 10.

Loadbalancers 14 and 18 are elements or devices that receive requests from one or more end users 12 and then distribute those requests to the next available server, gateway, or node. The available node may be any computer or device on a network that manages network resources, processes data, or receives information. For example, the server or node may be a selected one of a plurality of gateways 20 and 22 included within a gateway cluster. The loadbalancing decisions may be executed based on suitable algorithms, hardware, or software provided in loadbalancers 14 and 18. Loadbalancers 14 and 18 may also include suitable hardware and software for directing data or information in communication system 10. Additionally, the software or hardware within loadbalancers 14 and 18 may operate to build and store objects that correlate an IP address of end user 12 with a selected gateway 20 or 22. The correlation may be stored in a suitable table format that includes (in one example) four bytes allocated for the IP address of end user 12 being assigned to gateway 20 or 22. The IP address may be associated with version four (IPv4), version six (IPv6), or any other suitable communications protocol. Alternatively, the table may be replaced by any suitable element, guide, record, file, or object

operable to store the correlation between an IP address and an assigned gateway.

Loadbalancers 14 and 18 may also perform other suitable loadbalancing tasks, such as dividing the amount of work that an element has to do between two or more elements, ensuring more work gets done in the same amount of time and, in general, end users 12 can be served more quickly. Loadbalancers 14 and 18 may include any appropriate hardware, software, (or a combination of both) or any appropriate component, device, element, or object that suitably manages information traffic in a network environment. Additionally, any of the operations of gateways 20 and 22 may be assisted by loadbalancers 14 and 18 where appropriate and in accordance with particular needs.

In operation of an example embodiment, loadbalancers 14 and 18 may execute loadbalancing decisions for selected end users based on their IP addresses. From the perspective of loadbalancer 14 (#1), the end-user IP information may be checked and, based on that end-user IP information, end user 12 may be assigned to gateway (#1) 20. Thus, loadbalancer 14 may build an object to be included in a table that is stored in loadbalancer 14 such that subsequently received packets are directed to the same gateway. From the perspective of loadbalancer (#2) 18, the end-user IP address information within a packet may be checked. The first packet received by loadbalancer (#2) 18 (from the server perspective of the network, beyond gateways 20 and 22) may initiate the building of an object to be stored by loadbalancer (#2) 18. The object may again correlate an IP address to end user 12 such that subsequently received packets are

directed to the same gateway that handled the initial packet. In this sense, loadbalancing is being effectuated based solely on the end-user IP address (source IP or destination IP, depending on packet flow direction). Additional details relating to this process are provided below with reference to FIGURE 2.

Gateways 20 and 22 are networking elements that may perform any suitable operation in the network for one or more hosts that may each have a unique IP address. The hosts are end-users or end stations in the network that seek to communicate with each other or with another network entity or node. Gateways 20 and 22 may be part of a gateway cluster or single independent units. Gateways 20 and 22 may be a firewall, an intrusion detection system or element, a switch, a client service packet gateway (CSPG), a GPRS support node (GGSN), a packet data serving node (PDSN), a Layer-two tunneling protocol network server (LNS), or any other suitable device, component, object, node, or element operable to perform some function in a network environment.

Traffic may come from an appropriate Host A (e.g. end user 12), and loadbalancer (#1) 14 may then assign Host A to gateway (#1) 20. The function of each of loadbalancers 14 and 18 may be to ensure that traffic to and from Host A goes through the proper gateway (i.e. the one that has been assigned or stickied and provided in an object, which may be stored in a table within loadbalancer 14 or 18). Within many current existing network architectures, packet gateway functions may require that all IP traffic sourced from, or destined to, a particular host's IP address must pass through the same gateway. Gateways 20 and 22 may perform a variety of

per-host functions on the traffic of each host based on the identity of end user 12 (i.e. to what services or functions he/she is entitled or subscribed). Examples of these per-host operations are: network access/aggregation (e.g. PDSN, GGSN, LNS, dial, etc.), transparent content optimization transformation (e.g. WAP, optimization engines), security (e.g. secure sockets layer (SSL) termination, and firewalls), and personalized services and billing (e.g. CSPG). Increasing bandwidth and growing subscriber bases, as well as a demand for redundant gateway functions, have contributed to the need for loadbalancing gateways that perform such functions.

A CSPG is a client-aware device that may provide or offer some service or feature to end user 12. Such services (provided by a CSPG or either of gateways 20 and 22) may be based on an effective mapping between an end-user IP address of a given request packet and a user profile. These client-aware devices, reflected by gateways 20 and 22, may utilize the end-user IP address in providing services to end user 12. There are a number of reasons why a device or a component would want to identify the source or end user 12 associated with a communication session. For example, some devices may wish to identify end user 12 for authorization purposes. In another example, a device may wish to maintain user profiles to provide for accounting records (for example per-user accounting) or to provide for content billing information. Alternatively, a device or a component may use an identifier to provide for any other type of suitable client-aware service, tool, or feature according to particular needs of network components or equipment. Additional services may be related to areas such as

routing, permissions or access granting mechanisms, accounting, priority, billing information, firewalling, content filtering, or any other suitable parameters or policies where user-aware characteristics serve as a basis for network service implementation.

In other embodiments, gateways 20 and 22 may be provided and be used to terminate an end user's Layer-two point to point protocol (PPP) session and to actually *assign the IP addresses. In such a case, a simple* routing may be executed at loadbalancer (#1) 14, whereby one set of IP addresses may be directed to gateway (#1) 20 and another set of IP addresses may be directed to gateway (#2) 22. In other cases, gateways 20 and 22 may not be performing the address assignment, or the address pool may not be conducive to such segmenting amongst the gateways involved.

Servers 28 and 30 are locations in the network sought to be accessed by end user 12. Servers 28 and 30 may be considered as origin servers and include any suitable information or data to be accessed or communicated to end user 12. For example, server 28 could by yahoo.com, which may be operable to serve up web pages based on a request initiated by end user 12. Alternatively, servers 28 and 30 may be any other suitable element, location, or object sought to be accessed by end user 12.

FIGURE 2 is a simplified timing diagram illustrating an example interaction between multiple elements that may be included in communication system 10. The timing diagram may begin at step 1, where an end user initiates a request for a web page to be provided by yahoo.com. The request may be received by loadbalancer (#1) 14,

which may build an object that correlates the end user's IP address to the selected gateway. Loadbalancer (#1) 14 may then forward the packet/request on to gateway (#1) 20 at step 2.

In the example embodiment used for discussion and teaching purposes only, gateway (#1) 20 is a firewall. The firewall may generally seek to view all traffic to and from Host A and may also perform a number of other appropriate tasks, such as bandwidth monitoring for example. Such an operation could be tied into a billing system that may only allow a certain number of bytes to be delivered to an end user: denying service/connectivity beyond a certain amount of received data.

At step 3, gateway (#1) 20 may decide to forward the packet on to loadbalancer (#2) 18, which may forward the packet on to the appropriate server, at step 4. Loadbalancer (#2) 18 may view a MAC address (Layer-two) associated with the packet in order to glean its source. In the example provided, a TCP connection may be initiated by the end user. Thus, at step 5, server 28 may return the requested packet to loadbalancer (#2) 18, which has already built an appropriate object that correlates the end user's IP address to the selected gateway. Loadbalancer (#2) 18 may inspect the destination information in subsequently received packets in order to forward the packet to the appropriate gateway that is responsible for handling or processing information for the particular end user. Note that in other scenarios, loadbalancer (#2) 18 may opt to build an object when it receives any packet associated with Host A for the first time. Both loadbalancers 14 and 18 are capable of building objects at any appropriate time or

receiving information from an external element that provides information relating to locations to where incoming packets should be routed.

At step 6, loadbalancer (#2) 18 then forwards the packet to gateway (#1) 20, which may forward the packet on to loadbalancer (#1) 14 at step 7. Alternatively, step 7 may be avoided entirely in cases where the requested packet is delivered directly to the end user that requested the packet. Subsequently received packets may be properly directed by loadbalancers 14 and 18 based on end-user IP address information. The look-up is based on the end-user IP information the 'second time' this information is seen by the processing loadbalancer. A loadbalancing function is being provided in order to enforce the requirement that all traffic to and from a host needs to be routed in a certain manner. The object that was built by the participating loadbalancer is used for purposes of enforcement.

Some of the steps illustrated in FIGURE 2 may be changed or deleted where appropriate and additional steps may also be added to the flowchart. These changes may be based on specific communication architectures or particular interfacing arrangements and configurations of associated elements and do not depart from the scope or the teachings of the present invention.

Although the present invention has been described in detail with reference to particular communications, communication system 10 may be used for any clustered gateway architecture with requirements for load balancing all data traffic for a user to a single gateway. Any suitable communications that involve the use of end user 12 and network communications may benefit from the

teachings of the present invention, as elements such as gateways 20 and 22 have only been offered for purposes of example and should not be construed to limit the scope of the present invention in other potential environments. In addition, communication system 10 may be extended to any scenario for source or destination information associated with end user 12 that is capable of being discerned. This may be inclusive of Ethernet, token-ring, and other suitable protocols, architectures, configurations, and arrangements where appropriate and in accordance with particular needs.

Numerous other changes, substitutions, variations, alterations, and modifications may be ascertained by those skilled in the art and it is intended that the present invention encompass all such changes, substitutions, variations, alterations, and modifications as falling within the spirit and scope of the appended claims. Moreover, the present invention is not intended to be limited in any way by any statement in the specification that is not otherwise reflected in the appended claims.